

Science Target Assessment for Mars Rover Instrument Deployment

Liam Pedersen¹

¹*QSS Group, Inc. at NASA Ames Research Center, Moffett Field, California, USA (pedersen@email.arc.nasa.gov)*

Abstract

This paper introduces the system being developed at NASA Ames Research Center, intended for the Mars '09 Smart Lander, to robustly place sensors or tools against rocks in a single communications cycle. Science targets must be assessed prior to instrument placement in order to segment them from the background and determine where, if possible, to position the instrument. An initial result of this research effort is a novel Bayesian based method for segmenting rocks from the ground using 3D data.

1. Introduction

The ability to robustly place a manipulator mounted instrument against a science target in a single command cycle is essential for the proposed 2009 Mars Smart Lander rover mission (**Figure 1**). Without this level of autonomy the current science goals cannot be met in the time available for the mission.




Figure 1: Artist's conception of 2009 Mars Smart Lander [JPL]

The Sojourner rover (**Figure 2**), part of the 1997 Mars Pathfinder mission, required between three and five communications cycles between the rover and mission control to push an APX spectrometer against large rock targets. The operation was considerably simplified by the use of a compliant spring-loaded mechanism to push

the instrument against rocks. Furthermore, the lander was able to observe the rover and report to mission control on its progress. Mars Smart Lander will operate independently of a lander and will carry instruments requiring placement with centimeter accuracy from up to 10 m away.



Figure 2: Sojourner rover placing APX instrument against a rock target.[JPL]

Reliability and verifiability are the fundamental concerns for flight missions and the reasons why Sojourner had such limited autonomy. The rover could only execute rigid command sequences, the default response to unexpected behavior was to abort the sequence and wait for the next communications opportunity. The reasons for this is that these rigid sequences could be rigorously checked and verified by mission control prior to being uploaded to the vehicle, guaranteeing that a whole class of failure modes would not occur.

The NASA Ames Research Center has started an effort to develop robust single cycle instrument placement system on the Ames K9 rover platform (**Figure 3**) [1]. It will address the Mars '09 Smart Lander needs and support a demonstration of advanced autonomy (robust and intelligent execution, deliberative planning and autonomous science) in 2004.



Figure 3: *NASA Ames K9 rover platform.*

This paper introduces the NASA Ames single cycle instrument placement development effort describing the overall system architecture and component technologies. Central to this system are algorithms to assess potential target rocks to segment them from the ground underneath and determine where on them an instrument may be placed. An initial result is a novel 3D rock/ground segmentation algorithm developed for this purpose, and described herein.

1.1. Related work

There is a concurrent research effort [2] being at the Jet Propulsion Laboratory, using the Rocky 8 and FIDO rovers, that is also investigating robust instrument placement for Mars '09. Using accurate visual navigation techniques, they are able to accurately drive a rover up to a target and lower a camera onto it, stopping when the image is in focus. No contact is made with the target.

Rock/ground segmentation is also pertinent to the automated interpretation of rover science images. [3] demonstrate a different method to the method in this paper. It also used 3D data to find multiple rock targets in the viewing area. However, it assumes that the site is essentially planar, detecting rocks as small contiguous regions deviating from the plane. It is therefore unlikely to handle the cases likely to occur here, where a single target rock dominates the scene.

2. An architecture for robust instrument placement

Figure 5 (next page) indicates the sequence of actions required for a rover to robustly approach a science target and place an instrument against it using its manipulator arm. Mission controllers decide on a target and upload a flexible sequence with conditional

branches to the rover. This sequence includes the location of the target, information necessary to track it as the rover moves, and the actions and decisions needed to accomplish the task.

Specifying the instrument placement operation using a flexible execution sequence with conditional branches enables us to exploit ongoing research on robust execution [4][5] and the ground tools for generating such sequences [6]. The sequence constrains the range of possible behaviors and limits autonomous actions to well defined areas. Therefore, it can be verified and validated prior to uplink, increasing confidence that the rover will not get into a mission-endangering situation, such as low battery level.

Subsequent to target designation and sequence upload, the rover must navigate to a location in from of the target while avoiding obstacles and keeping track of the target's location, perhaps using the visual servoing algorithm of [7].




Figure 4: *NASA Ames K9 rover stereo navigation and science camera pairs.*

Upon getting close to the target the rover must determine its 3D shape (using high resolution stereo cameras, **Figure 4**) and precise location relative to the robot. This information is used to assess the target: segmenting the rock from the background soil, finding which areas of the rock are consistent with the instrument and science requirements and which of these are in the manipulator workspace.

If part of the rock target, consistent with the instrument and science requirements, can be reached, then the desired instrument pose must be calculated and the arm deployed, using contact or force sensors to determine when the instrument is in contact with the rock after the final few centimeters of motion (**Figure 6**). Otherwise, a re-plan is necessary, either by mission control (who would upload a new sequence to move the rover after the next communications cycle), or possibly by an onboard planner.

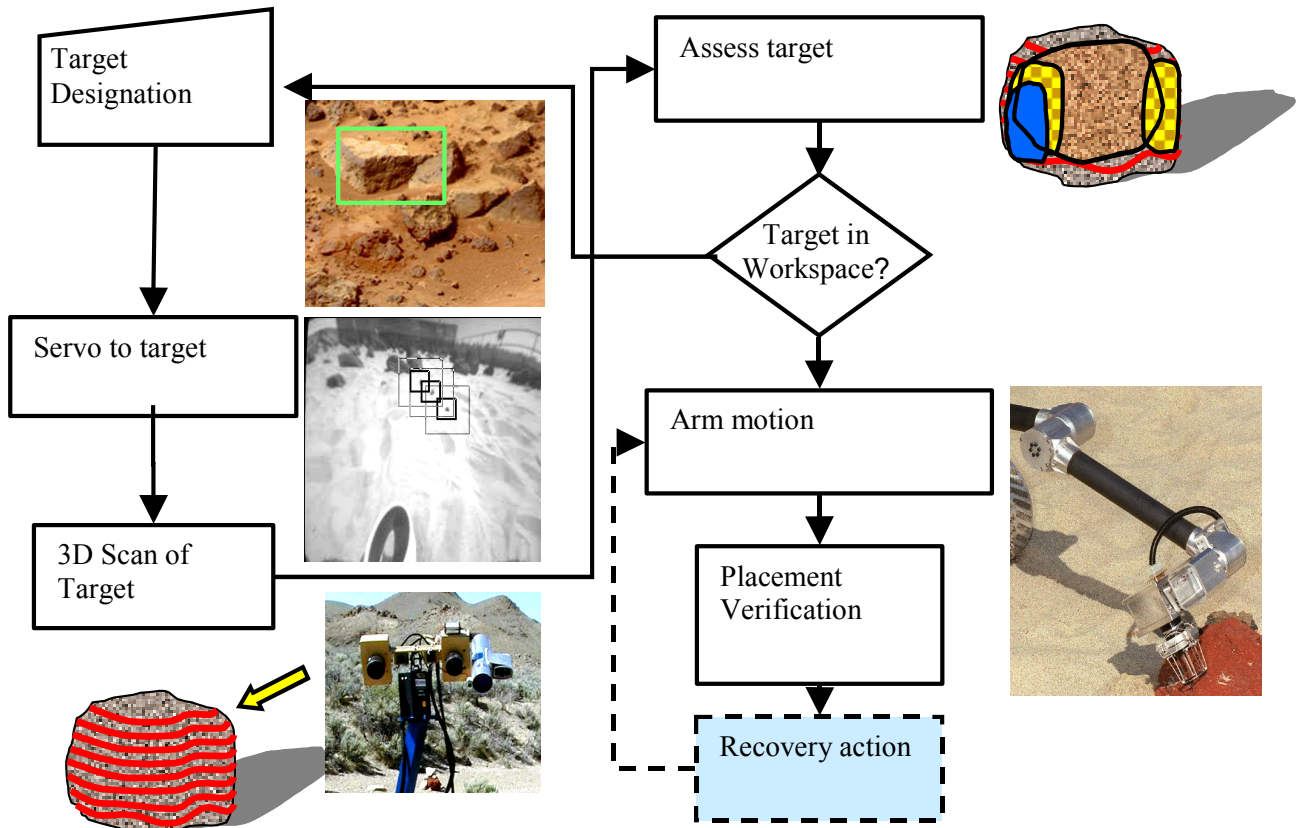


Figure 5: Sequence of actions for a rover to approach a science target from up to 10m distance and use a manipulator arm to place an instrument against the target in a manner consistent with the instrument operating criteria and the science needs.



Figure 6: K9 Rover 5 DOF manipulator arm with the CHAMP microscopic imager attached.

3. 3D Rock and ground segmentation

Assuming the rover has scanned the rock target and its background to obtain a 3D dot cloud $\mathbf{X}=\{\mathbf{x}_i|i=1..N\}$, the rock ground segmentation problem can be restated as that of finding their most likely rock/ground classifications $\mathbf{C}=\{\mathbf{c}_i|i=1..N\}$, integrated over the

distribution of possible values for the scene model parameters \mathbf{M} describing the rocks and the ground and the prior probabilities \mathbf{P} of each. Both \mathbf{P} and \mathbf{M} are unknown “hidden” or “nuisance” variables.

This is the classic problem of estimating the modes of $P(\mathbf{C}|\mathbf{X})$, integrating over the hidden variables, and solved using the *Expectation-Maximization* (EM) algorithm [8][9], which in this case can be roughly stated as follows:

1. Pick initial estimates for \mathbf{M} and \mathbf{P} .
2. E-step: Determine maximum likelihood values for \mathbf{C} given the estimates of \mathbf{M} and \mathbf{P} , and the measurements \mathbf{X} . In this case the individual maximum likelihood classifications \mathbf{c}_i can be computed independently by maximizing $P(\mathbf{c}_i)p(\mathbf{x}_i|\hat{\mathbf{M}},\mathbf{c}_i)$
3. M-step: compute new the “best-fit” estimates for \mathbf{M} and \mathbf{P} from \mathbf{X} and max likelihood \mathbf{C} based on the previous parameter estimates. Iterate until convergence.

In the special case that $p(x_i | M, c_i) \propto \exp(D(x_i, c_i, M))$, where $D(\cdot)$ is a distance metric between the point x_i and its corresponding “class center”, the algorithm reduces to a variation of the k-means clustering algorithm:

1. Pick initial estimates for \mathbf{M} and \mathbf{P} .
2. Classify each \mathbf{x}_i according to whether it is closer to the rock class or the ground class as defined by the latest estimates for the model parameters \mathbf{M} and class prior probabilities \mathbf{P} .
3. Re-estimate \mathbf{M} and \mathbf{P} given \mathbf{C} and iterate until convergence.

In step 3 above, the parameters \mathbf{M} for a given class are estimated by minimizing the sum of square errors as defined by the class distance metric.

3.1. Rock and ground scene model

In order to apply the previous algorithm it is necessary to assume a scene model describing rocks, the ground and the distance between a 3D point and a given rock or the ground.

Within the limited area viewable in front of the robot, the ground can be usefully modeled as a plane, described by a unit normal vector \mathbf{n} and the perpendicular distance to the origin in the robot coordinate frame (**Figure 7**). Rocks are modeled as hemispheres, described by their position and radius, and constrained to always point towards the viewing camera. This is because the camera pair used to obtain range information can only see the front side of a rock and a sphere is therefore an inadequate model.

Distances between data points and the ground plane are trivially computed. To compute the distance between a data point and a hemisphere ignore the back-plane and consider only the curved surface.

Given a set of data points that are known to belong to the ground, the best-fit ground plane is easily determined by computing the eigenvector corresponding to the smallest eigenvalue. This is the unit normal vector to the plane. Computing the best-fit hemisphere to a set of points is best done by numerically minimizing the sum of square errors using an algorithm such as the Nelder-Mead simplex (direct search) method [10].




Figure 7: *Rock and ground scene model*

3.2. Experimental rock/ground segmentation results

Stereo camera (NavCams) images were acquired of a rock pile (**Figure 8**) in front of the K9 rover and a range map computed using the Ames stereo pipeline algorithm [11], and then transformed into a 3D dot cloud.




Figure 8: *Rock scene in front of the K9 robot.*

This dot cloud was then segmented into rock and ground plane pixels using the k-means based algorithm above, assuming a single rock object in the scene. **Figure 9** shows the resulting class assignments (segmentation) for the 3D data points. The rock and ground plane points are clearly distinguished, in spite of the considerable noise and non-spherical nature of the rock.




Figure 9: 3D point cloud derived from stereo range data of the scene in **Figure 8** and segmented into points belonging to the rock pile (red) and points belonging to the ground plane (green).




Figure 10: Top view of segmented dot cloud in **Figure 9**.

The derived scene model based on maximum likelihood “nuisance” model parameters (rock position and radius, ground plane orientation and distance from origin) is shown in **Figure 11**.




Figure 11: Derived scene model from the segmentation of the 3D data from **Figure 8**.

4. Conclusions and future work

The capability to place sophisticated instruments against science targets in one command cycle from 10m away is essential for forthcoming planetary exploration rovers.

The proposed overall instrument placement architecture rover architecture based on combining capabilities together with robust execution engines and ground based sequence-generating tools addresses the robustness and verification issues faced by flight missions. Comprehensive and systematic rover field trials as the system is developed over the next few years are planned to validate this claim.

A necessary capability for instrument placement is the segmentation of 3D data points generated by a target rock pile from those generated from the ground. The algorithm presented here accomplishes this, assuming only that the ground underlying the target rock is relatively flat, and one compact rock pile is present. No arbitrary parameters (other than bounds on the workspace) are used, and supervised training is unnecessary. Furthermore, the algorithm is robust to noise, missing data, sparsely sampled or occluded regions. It can be run on low-resolution data, minimizing execution time without compromising performance.

Improved robustness could be attained by incorporating color and texture into the scene model. This algorithm is very amenable to such changes, requiring only a refinement of the scene model.

The main drawback of the clustering algorithm demonstrated here is the difficulty of robustly extending it to detecting an arbitrary number of rocks in the scene. This can be done by the algorithm in [3], but this is more complex and conceivably would fail when the scene is dominated by a single, large rock pile, as is likely in this scenario.

Future work in this project will concentrate on the development of an end-to-end rover instrument placement system. It will likely incorporate many different technologies. The biggest challenges will be linking them together and understanding the behavior of the complete system.

5. Acknowledgements

This work is an ongoing project in the Intelligent Robotics Group in Code IC at NASA Ames Research Center, and is supported by NASA’s Intelligent System’s program. Infrastructure for this project is being developed with support from NASA’s Mars Technology Program.

6. References

- [1] J. Bresina, M. Bualat, L. Edwards, R. Washington, A. Wright, "K9 operations in May '00 dual-rover field experiment", *i-SAIRAS*, 2001.
- [2] T. Huntsberger, H. Aghazarian, Y. Cheng, E.T. Baumgartner, E. Tunstel, C. Leger, A. Trebi-Ollennu and P. Schenker, "Rover Autonomy for Long Range Navigation and Science Data Acquisition on Planetary Surfaces", *IEEE International Conference on Robotics and Automation*, 2002.
- [3] V. Gor, R. Castano, R. Manduchi, R.C. Anderson, and E. Mjolsness, "Autonomous Rock Detections for Mars Terrain", *AIAA*, 2001.
- [4] J. Bresina, K. Golden, D.E. Smith, R. Washington, "Increased Flexibility and Robustness for Mars Rovers", *i-SAIRAS* 1999.
- [5] N. Muscettola, G.A. Dorais, R. Levinson and C. Plaunt, "IDEA: Unifying Planning and Execution for Real-Time Autonomy", *i-SAIRAS* 2001.
- [6] Kanna Rajan, Mark Shirley, Will Taylor and Bob Kanefsky, "Ground Tools for the 21st Century" *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, 2000.
- [7] Mark W. Maimone, Issa A. Nesnas, Hari Das, "Autonomous Rock Tracking and Acquisition from a Mars Rover", *i-SAIRAS*, 1999.
- [8] C. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press Inc., New York, 1995
- [9] Gelman, J.B. Carlin, H.S. Stern, and D.B. Rubin, *Bayesian Data Analysis*, Chapman & Hall, 1995.
- [10] W.H. Press, W.T. Vetterling, S.A. Teukolsky and B.P. Flannery, *Numerical Recipes in C*, Cambridge University Press, 1992.
- [11] Stoker, C. et al., "Analyzing Pathfinder Data Using Virtual Reality and Super-resolved Imaging", *Journal of Geophysical Research*, vol. 104, no. E4, pp.8889-8906, April 25, 1999.